

Multi-Dimensional Balanced Graph Partitioning via Projected Gradient Descent*

Dmitrii Avdiukhin
Indiana University
Bloomington, IN
davdyukh@iu.edu

Sergey Pupyrev
Facebook
Menlo Park, CA
spupyrev@gmail.com

Grigory Yaroslavtsev
Indiana University
Bloomington, IN
grigory@grigory.us

ABSTRACT

Motivated by performance optimization of large-scale graph processing systems that distribute the graph across multiple machines, we consider the balanced graph partitioning problem. Compared to most of the previous work, we study the *multi-dimensional* variant in which balance according to multiple weight functions is required. As we demonstrate by experimental evaluation, such multi-dimensional balance is essential for achieving performance improvements for typical distributed graph processing workloads.

We propose a new scalable technique for the multi-dimensional balanced graph partitioning problem. It is based on applying randomized projected gradient descent to a non-convex continuous relaxation of the objective. We show how to implement the new algorithm efficiently in both theory and practice utilizing various approaches for the projection step. Experiments with large-scale graphs containing up to hundreds of billions of edges indicate that our algorithm has superior performance compared to the state of the art.

PVLDB Reference Format:

Dmitrii Avdiukhin, Sergey Pupyrev, Grigory Yaroslavtsev. Multi-Dimensional Balanced Graph Partitioning via Projected Gradient Descent. *PVLDB*, 12(8): 906 - 919, 2019.
DOI: <https://doi.org/10.14778/3324301.3324307>

1. INTRODUCTION

Distributed graph processing systems have been widely adopted in recent years to enable analysis and knowledge extraction from large-scale graphs. Systems such as Giraph [6], GraphX [19], GraphLab [30], and PowerGraph [18] allow users to use a vertex-centric model for applications which can be executed on a cluster of worker nodes. In this setting, each worker node operates on a subset of the input graph and communicates with other workers by sending messages.

*Research supported by NSF award CCF-1657477 and Facebook Faculty Research Award. Full version: <http://arxiv.org/abs/1902.03522>.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 12, No. 8

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3324301.3324307>

The process of splitting the input graph into these subsets, also known as graph partitioning, is essential for optimizing performance of such systems [3, 18, 20, 43].

Created partitions have a significant impact on the communication between different workers and the resource usage of individual workers. In order to maximize the processing speed, the partitions should largely be independent to minimize communication. At the same time, computation executed on each partition should take approximately the same amount of processing time, as the overall performance depends on the slowest worker. These constraints give rise to the BALANCED GRAPH PARTITIONING problem whose goal is to divide the vertices of a graph into a given number of (approximately) equal size components while minimizing the resulting edge cut. BALANCED GRAPH PARTITIONING is a classic and thoroughly studied problem from both theoretical and practical points of view [9, 12]. In the context of distributed graph processing, the problem is typically studied in two variants.

In the *vertex partitioning* model, each worker machine is assigned an equal number of vertices with the goal of minimizing the number of cross-machine edges. Since messages are usually sent between adjacent vertices, storing tightly connected subgraphs on the same worker can reduce communication and hence running times of jobs. It has however been observed that this strategy does not lead to equally loaded partitions for real-world graphs with power law degree distribution [18]. Graph partitioning algorithms tend to collocate high-degree vertices and corresponding partitions take much longer to process, resulting in longer execution time overall.

The *edge partitioning* model has been suggested to alleviate the above imbalance problem [18, 29]. In this model the goal is to partition the graph so that the number of edges in every component is the same, while the number of incident edges across different components is minimized. Good partitions according to this model typically result in better balance across workers and reduced computation time in comparison to the trivial hash-based assignment of vertices to worker machines. However, edge-based graph partitioning can still result in performance regressions [3, 40].

To analyze the source of regressions, we performed a simple experiment of running a Page Rank algorithm implemented on top of Giraph utilizing various graph partitioning methods. Figure 1 illustrates the histograms of running times for individual workers processing a graph with 800M vertices and 80B edges. As discussed above, partitions according to the *vertex partitioning* model suffer from unequal distribu-




Figure 1: The running time of an iteration of Page Rank on a Giraph cluster of 16 worker machines using various graph partitioning strategies. The numbers indicate the average percentage of local (uncut) edges per worker, which is proportional to the number of local messages for the distributed graph processing workload. Vertex-edge partitioning achieves approximately 12% iteration time improvement compared to hash.

tion of edges across workers. A single overloaded partition contains 1.92x more edges than an average one, which results in 1.5x longer execution time. We also observe a high correlation ($\rho = 0.79$) between the number of edges assigned to a partition and the corresponding processing time in this experiment. Partitioning according to the *edge partitioning* model yields a 1.08x running time improvement over the baseline, though there is still a noticeable imbalance between the fastest and the slowest worker machines. This can be explained by uneven distribution of vertices among workers. Machines with more vertices have higher operational overhead such as serialization of sent messages whose number is proportional to the number of vertices on a worker. Here we observe an 1.33x imbalance in the number of vertices and a moderate correlation ($\rho = 0.62$) between the running time and the vertex count on the workers.

In order to mitigate the issues described above we introduce a new strategy, *vertex-edge partitioning*, which is designed to balance the number of vertices and edges across workers *simultaneously*. As shown in Figure 1, this is done at a cost of a lower edge locality (percentage of edges with both endpoints on the same machine), and thus, higher communication volume. The resulting assignment results in a 1.12x speedup over the hash-based model. Motivated by the above experiment and a number of earlier studies [3, 20, 40, 43], we formalize a new model for graph partitioning, which is suitable for real-world distributed graph processing systems.

We now formally describe the model in the most general setting which allows one to require balance according to d different unrelated weight functions. Let $G(V, E)$ be a graph with d vertex weight functions $w^{(1)}, \dots, w^{(d)} : V \rightarrow \mathbb{R}^+$, each assigning a positive weight to every vertex in the graph. Let $w^{(j)}(V) = \sum_{v \in V} w^{(j)}(v)$ be the sum of weights of all vertices in the graph according to the j -th weight function. Given an integer k and a parameter $\varepsilon > 0$, the goal is to partition the vertex set V into k sets V_1, \dots, V_k such that, for each weight function $w^{(j)}$ and each part V_i , the sum of weights in V_i is approximately the same and close to the average, that is, $\sum_{v \in V_i} w^{(j)}(v) = (1 \pm \varepsilon) \frac{w^{(j)}(V)}{k}$. We call such partitions ε -balanced. Finally, among all such ε -balanced partitions the goal is to find one that maximizes the number of edges whose both endpoints are contained within some component of the partition and hence minimizes the size of the cut. This problem is referred to

as MULTI-DIMENSIONAL BALANCED GRAPH PARTITIONING (MDBGP).

The simplest example of MDBGP is the classic *balanced graph partitioning problem*, which is equivalent to the vertex partitioning strategy described above and can be expressed using a single weight function $w^{(1)}(v) = 1$. Since $w^{(1)}(V) = |V|$ this requires that we maximize edge locality while ensuring that $|V_i| \approx \frac{|V|}{k}$. Using two weight functions, $w^{(1)}(v) = 1$ and $w^{(2)}(v) = \text{deg}(v)$, corresponds to requiring balance on the number of vertices and edges in the parts of the partition and hence corresponds to the vertex-edge partitioning approach mentioned above. Indeed, $w^{(2)}(V) = 2|E|$ and hence in addition to balance on the number of vertices this requires that $\sum_{v \in V_i} \text{deg}(v) \approx \frac{2|E|}{k}$. However, the model is not restricted to vertex- and edge-balance (as in the aforementioned *vertex-edge partitioning*) but can take arbitrary user-specified weights. In particular, when partitioning the vertices of the graph between the workers for load balancing, various weights modeling expected vertex activity can be used (historical data on individual vertex load, proxy values for the load such as PageRank, etc).

While a large body of work exists offering practical solutions for the one-dimensional version of the problem [7, 12, 13, 14, 22, 23, 33, 41, 42], as well as on theoretical foundations of graph partitioning [4, 26, 32], literature on principled and scalable approaches for the multi-dimensional case is quite sparse [24, 35, 36, 37]. In particular, if the weight functions are unrelated to each other, one can easily construct examples when no feasible solution exists that satisfies all balance constraints even for two weight functions. However, it is empirically observed that instances coming from applications often allow balanced solutions for several weight functions of interest simultaneously. For classical local search based algorithms such as [25], handling of multiple unrelated weight functions is challenging since imposing one balance constraint might violate another; hence, finding a good local move becomes computationally hard. We overcome this difficulty by using a continuous relaxation of the problem. In order to obtain an integral solution, we apply randomized rounding which preserves balance with high probability.

1.1 Our Contributions

We present a scalable algorithmic framework for the problem of balanced partitioning of large graphs according to

multiple user-specified weight functions while maximizing the number of edges inside the resulting components. Our framework consists of applying the *projected gradient descent* on a standard relaxation with a suitably chosen projection method. The relaxation is to maximize a non-convex quadratic function $f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T A \mathbf{x}$ for $\mathbf{x} \in \mathbb{R}^n$, where A is the adjacency matrix, subject to a constraint $\mathbf{x} \in K$ for a certain convex body K defined by the weight functions. Section 2 provides the exact description of the relaxation. Note that the gradient descent step only uses a matrix-vector multiplication since $\nabla f = A\mathbf{x}$. Thus, the algorithm allows a straightforward distributed implementation.

While applying projected gradient descent to solve non-convex optimization problems subject to convex constraints is a well-studied approach in non-linear optimization (Section 2.3, [8]) and machine learning (Section 6.6, [21]), one has to overcome two technical challenges to make it applicable to the multi-dimensional graph partitioning problem: 1) projection step is computationally expensive, 2) existence of points with small gradient (saddle points) slows down convergence.

We show how to address the first challenge by designing efficient projection step algorithms tailored to the standard relaxation of MDBGP. While convergence to the projection point can be achieved using various alternating projections methods [15], for $d \leq 2$ we give one-shot exact solutions with almost linear running time.

THEOREM 1.1. *Running time of the projected gradient descent step is $\mathcal{O}(|E| + |V| \log^{d-1} |V|)$ for $d \leq 2$ and scales as $\mathcal{O}(|E|/m + |V| \log^{d-1} |V|)$ when distributed over m machines.*

In order to address the second challenge, we use small perturbations to get out of saddle points, where the perturbation vectors are sampled from a scaled n -dimensional Gaussian distribution. We refer to the resulting algorithm as **Gradient Descent (GD)**, see Algorithm 1. Convergence analysis of GD remains an open problem. While noisy gradient descent is known to have fast convergence to a local optimum for non-convex optimization subject to equality constraints, if inequality constraints are allowed, convergence analysis is unknown [16].

Our experiments show that GD scales to graphs with up to several billions of vertices and several hundreds of billions of edges. We perform a comparison of various graph partitioning strategies for optimizing several real-world Giraph workloads. The results demonstrate that multi-dimensional balancing is a suitable objective for achieving performance improvements, leading to speedups in the order of 10% – 30% over the state-of-the-art one-dimensional strategies. Compared to existing scalable graph partitioners, such as Social Hash Partitioner [22], Spinner [33], and Balanced Label Propagation [34, 42], the algorithm is conceptually simple and obtains close-to-perfect balance across multiple dimensions.

1.2 Previous Work

While one-dimensional balanced graph partitioning has been studied extensively and a number of tools exist [7, 13, 14, 22, 23, 33, 41, 42] (see also surveys by Bichot and Siarry [9] and by Buluç et al. [12]), to the best of our knowledge none of the practical algorithms for this problem have been previously based on running gradient descent on a continuous relaxation. Existing approaches are inherently discrete and are based on combinations of various discrete algorithms: greedy heuristics

(METIS [23], Fennel [41]), branch-and-bound [13], label propagation and local search (Balanced Label Propagation [42], Social Hash Partitioner [22], Spinner [33]), as well as hybrid approaches (linear embedding method combined with various optimizations [7]). Due to the combinatorial nature of these algorithms, their generalizations to the multi-dimensional case appear to be non-straightforward without substantial losses in performance, while our continuous relaxation handles multiple balance constraints uniformly. Compared to the one-dimensional version, existing literature on the multi-dimensional version is rather sparse [24, 35, 36, 37] and the main publicly available tool for the problem is currently METIS [24, 37].

Vast literature exists on optimization of non-convex functions and the interest in this topic lately has been particularly high. However, in the constrained case when the optimization has to be performed over a convex body, fairly little is known; see classic optimization literature [8, 11, 44]. Recent results on the non-convex optimization problem subject to convex constraints and its special cases include [5, 16, 17, 21, 39]. Closest to our work in terms of techniques is [27] who use projected gradient method to solve convex programs involving the max-norm and show how to solve large semidefinite programming relaxations of MAX-CUT. Their results are quite different from ours as we consider a balanced version of graph partitioning and expect our algorithms to be scalable; the largest instances handled by [27] have $|V| = 20K$ and $|E| = 40K$. Since we require that our algorithms scale to graphs with billions of edges, using existing general purpose software for constrained quadratic programming is infeasible.

2. PROJECTED GRADIENT DESCENT

For an integer $t \geq 1$, we use notation $[t]$ to denote the set $\{1, \dots, t\}$. The weighted d -dimensional balanced graph partitioning problem is defined by a collection of d weight functions $w^{(1)}, \dots, w^{(d)}$, where $w^{(j)}: V \rightarrow \mathbb{R}^+$. For a set $S \subseteq V$ we use notation $w^{(j)}(S) \equiv \sum_{v \in S} w_v^{(j)}$.

DEFINITION 2.1 (MDBGP). *For a graph $G(V, E)$, the MULTI-DIMENSIONAL ε -BALANCED GRAPH k -PARTITIONING problem is to partition the vertex set V into k sets V_1, \dots, V_k such that for each $j \in [d]$, it holds that $w^{(j)}(V_i) = (1 \pm \varepsilon) \frac{w^{(j)}(V)}{k}$ for all $i \in [k]$. Among all such partitions the goal is to find one that maximizes the number of edges whose both endpoints are within some part.*

In this paper we focus on the 2-partitioning problem; for the general variant of k -partitioning, we apply the algorithm recursively. For $k = 2$ MDBGP is equivalent to the following integer quadratic program:

$$\begin{aligned} \text{Maximize:} & \quad \frac{1}{2} \sum_{(i_1, i_2) \in E} (x_{i_1} x_{i_2} + 1) \\ \text{Subject to:} & \quad \left| \sum_{i=1}^n w_i^{(j)} x_i \right| \leq \varepsilon \sum_{i=1}^n w_i^{(j)} \quad \forall j \in [d] \\ & \quad x_i \in \{-1, 1\} \quad \forall i \in V \end{aligned}$$

The interpretation of x_i variables is that if $x_i = 1$ then $i \in V_1$ and if $x_i = -1$ then $i \in V_2$. The objective is then the same as in MDBGP and counts the number of edges

whose both endpoints are contained in some part of the partition. Indeed, an edge (i_1, i_2) makes a contribution of 1 to the objective when $x_{i_1} = x_{i_2}$ (and hence $x_{i_1}x_{i_2} = 1$) and 0, otherwise (since $x_{i_1}x_{i_2} = -1$). The constraints are equivalent to $-\varepsilon w^{(j)}(V) \leq w^{(j)}(V_1) - w^{(j)}(V_2) \leq \varepsilon w^{(j)}(V)$. Adding or subtracting $w^{(j)}(V)$ to both sides and dividing by 2, we have $w^{(j)}(V_i) = (1 \pm \varepsilon) \frac{w^{(j)}(V)}{2}$ as required in MDBGP.

After dropping the additive term, the objective can be expressed as $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T A \mathbf{x}$ and has gradient $\nabla f(\mathbf{x}) = A \mathbf{x}$ and Hessian $\nabla^2 f = A$. Finally, we use a continuous relaxation of the above problem where we replace the integrality constraints with $x_i \in [-1, 1]$ for all $i \in V$. A solution to this continuous relaxation can be converted into an integral solution using randomized rounding. Using independent random variables X_i for each vertex such that $\Pr[X_i = 1] = \frac{1+x_i}{2}$ and $\Pr[X_i = -1] = \frac{1-x_i}{2}$ the expected value of the objective on the rounded solution $(X_1, \dots, X_{|V|})$ is the same as on the initial fractional solution $(x_1, \dots, x_{|V|})$ while all balance constraints are still approximately preserved with high probability by concentration bounds.

2.1 Overview

We propose the following algorithm for the multi-dimensional balanced graph partitioning problem based on the continuous relaxation described above. The algorithm is referred to as **Gradient Descent (GD)**, see Algorithm 1. It computes a sequence of vectors $\{\mathbf{x}^{(t)}\}$, where $x_i^{(t)} \in [-1, 1]$ for all $i \in V$ and t . Here $\mathbf{x}^{(0)}$ is initialized with zero vector, and $\mathbf{x}^{(t+1)}$ is computed by applying projected gradient descent iteration to $\mathbf{x}^{(t)}$. Each iteration consists of three steps.

Step 1: Adding noise. We add Gaussian noise to $\mathbf{x}^{(t)}$ obtaining a noisy vector $\mathbf{z}^{(t)}$. The noise is drawn from the n -dimensional Gaussian distribution $N_n(0, \eta_t)$ with zero mean and variance η_t in each coordinate. The addition of noise to $\mathbf{x}^{(t)}$ allows to escape from saddle points, e.g. $\mathbf{x}^{(0)} = 0$.

Step 2: Gradient descent. We obtain $\mathbf{y}^{(t)}$ from the noisy vector $\mathbf{z}^{(t)}$ via a gradient descent step with step size γ_t . Note that the gradient at $\mathbf{z}^{(t)}$ is given as $A \mathbf{z}^{(t)}$ hence this step can be expressed as $\mathbf{y}^{(t)} = (I + \gamma_t A) \mathbf{z}^{(t)}$.

Step 3: Projection. The resulting vector $\mathbf{y}^{(t)}$ is then projected on the feasible space $\mathcal{B}_\infty \cap \bigcap_{j=1}^d \mathcal{S}_\varepsilon^j$, where:

$$\mathcal{B}_\infty = \{ \mathbf{x} \in \mathbb{R}^n \mid \forall i: x_i \in [-1, 1] \}$$

$$\mathcal{S}_\varepsilon^j = \{ \mathbf{x} \in \mathbb{R}^n \mid \left| \sum_{i=1}^n w_i^{(j)} x_i \right| \leq \varepsilon \sum_{i=1}^n w_i^{(j)} \} \text{ for } j \in [d],$$

that is, \mathcal{B}_∞ satisfies that $\|\mathbf{x}\|_\infty \leq 1$ and $\mathcal{S}_\varepsilon^j$ corresponds to the constraints imposed by the balance of weights according to the j -th weight function.

The final solution is obtained by rounding last $\mathbf{x}^{(t)}$: each vertex i is assigned to part V_1 with probability $\frac{x_i^{(t)} + 1}{2}$. Note that this ensures that the expected number of edges whose endpoints belong to the same part after this rounding is given as $\frac{1}{2} \sum_{(i_1, i_2) \in E} (x_{i_1}^{(t)} x_{i_2}^{(t)} + 1)$.

The algorithm uses parameters η_t, γ_t , and I , where t is the iteration index. Here η_t controls the magnitude of noise, γ_t is the step size, and I is the number of iterations. We discuss the selection of parameters in the experimental Section 4.

2.2 Projection

Algorithm 1: GD (d -Dimensional Balanced Graph 2-Partitioning via Randomized Projected Gradient Descent)

input : Graph $G(V, E)$, $\varepsilon \in [0, 1]$, weight functions $w_1, \dots, w_d: V \rightarrow \mathbb{R}^+$
parameters: $I, \{\eta_t\}_{t=0}^{I-1}, \{\gamma_t\}_{t=0}^{I-1}$
output : ε -balanced partition w.r.t $w^{(1)}, \dots, w^{(d)}$ of V into (V_1, V_2)

- 1 $K = \mathcal{B}_\infty \cap \bigcap_{j=1}^d \mathcal{S}_\varepsilon^j$
- 2 $\mathbf{x}^{(0)} = \mathbf{0}$;
- 3 **for** $t = 0$ **to** $I - 1$ **do**
- 4 $\mathbf{z}^{(t)} = \mathbf{x}^{(t)} + N_n(0, \eta_t)$; // Noise addition step
- 5 $\mathbf{y}^{(t+1)} = (I + \gamma_t A) \mathbf{z}^{(t)}$; // Gradient descent step
- 6 $\mathbf{x}^{(t+1)} = \operatorname{argmin}_{\mathbf{x} \in K} \|\mathbf{y}^{(t+1)} - \mathbf{x}\|_2$; // Projection step
- 7 $V_1 = V_2 = \emptyset$; // Randomized rounding
- 8 **for each** $i \in V$ **do**
- 9 With probability $\frac{x_i^{(t)} + 1}{2}$, let $V_1 = V_1 \cup \{i\}$;
 otherwise, $V_2 = V_2 \cup \{i\}$;

In the projection step of GD (Line 6) we need to find $\operatorname{argmin}_{\mathbf{x} \in K} \|\mathbf{y}^{(t+1)} - \mathbf{x}\|_2$, where $K = \mathcal{B}_\infty \cap \bigcap_{j=1}^d \mathcal{S}_\varepsilon^j$. Denoting $\mathbf{y}^{(t+1)}$ as \mathbf{y} we formulate this step as an optimization problem:

$$\begin{aligned} \text{Minimize:} \quad & f(\mathbf{x}) = \|\mathbf{x} - \mathbf{y}\|_2^2 \\ \text{Subject to:} \quad & g_i = x_i^2 - 1 \leq 0 \quad \forall i \in [n] \\ & h_+^{(j)} = \sum_{i=1}^n w_i^{(j)} x_i - \varepsilon \leq 0 \quad \forall j \in [d] \\ & h_-^{(j)} = -\sum_{i=1}^n w_i^{(j)} x_i - \varepsilon \leq 0 \quad \forall j \in [d] \end{aligned}$$

The optimum solution to the optimization problem has to satisfy KKT conditions:

$$\begin{aligned} \text{Stationarity:} \\ \mathbf{y} - \mathbf{x} &= \sum_{i=1}^n \mu_i x_i \mathbf{e}_i + \sum_{j=1}^d (\mu_+^{(j)} - \mu_-^{(j)}) \sum_{i=1}^n w_i^{(j)} \mathbf{e}_i \\ \text{Complementary slackness 1:} \\ \mu_i (x_i^2 - 1) &= 0, \quad \forall i \in [n] \\ \text{Complementary slackness 2:} \\ \mu_+^{(j)} \left(\sum_{i=1}^n w_i^{(j)} x_i - \varepsilon \right) &= 0, \quad \forall j \in [d] \\ \mu_-^{(j)} \left(-\sum_{i=1}^n w_i^{(j)} x_i - \varepsilon \right) &= 0, \quad \forall j \in [d] \end{aligned}$$

Here $\mu_i, \mu_+^{(j)}, \mu_-^{(j)} \geq 0$ are the dual variables and \mathbf{e}_i is the i -th standard unit vector. It is a standard fact (see [11], Chapter 5.5.3) that for convex optimization subject to linear constraints Stationarity, Complementary slackness and Primal/Dual feasibility are necessary and sufficient conditions

for the optimum solution. Thus we just focus on satisfying these conditions below.

Let $\gamma_i = \sum_{j=1}^d (\mu_+^{(j)} - \mu_-^{(j)}) w_i^{(j)}$. Then by Stationarity for each i we have $y_i - x_i = \mu_i x_i + \gamma_i$. Consider the following three cases:

Case 1. ($y_i > 1 + \gamma_i$). If $\mu_i = 0$ then by Stationarity $x_i = y_i - \gamma_i > 1$ which violates primal feasibility conditions. Therefore $\mu_i > 0$ and $x_i^2 = 1$ by Complementary slackness 1. Among the two roots $x_i = 1$ and $x_i = -1$ the second root can be ruled out and hence $x_i = 1$. Indeed, if $x_i = -1$ then by Stationarity $y_i + 1 = -\mu_i + \gamma_i$ which contradicts $\mu_i > 0$ and $y_i > 1 + \gamma_i$.

Case 2. ($y_i < -1 + \gamma_i$). This case is symmetric to the previous one and thus $x_i = -1$ in this case.

Case 3. $y_i \in [-1 + \gamma_i, 1 + \gamma_i]$. First we show that $\mu_i = 0$. Indeed, assume that $\mu_i > 0$. Then $x_i = \pm 1$ by Complementary slackness 1. Both cases lead to contradiction:

1. ($x_i = 1$). By Stationarity $y_i - 1 = \mu_i + \gamma_i$ which contradicts with $y_i \leq 1 + \gamma_i$ and $\mu_i > 0$.
2. ($x_i = -1$). Similarly to the above by Stationarity we have $y_i + 1 = -\mu_i + \gamma_i$ which is a contradiction with $y_i \geq -1 + \gamma_i$ and $\mu_i > 0$.

Therefore in this case we have $\mu_i = 0$ and hence by Stationarity $x_i = y_i - \gamma_i$.

Let $\lambda_j = \mu_+^{(j)} - \mu_-^{(j)}$ and assume that these values are known to the algorithm. For $z \in \mathbb{R}$ we use notation $[z] = \min(1, \max(-1, z))$ for the truncated linear function. Using the analysis above the projection step is simply $x_i = [y_i - \sum_{j=1}^d \lambda_j w_i^{(j)}]$. It remains to show how to find $\{\lambda_j\}$.

Note that from Complementary slackness 2 it follows that either $\mu_+^{(j)} = 0$ or $\mu_-^{(j)} = 0$ since both of these values being positive leads to a contradiction. This leads to three cases: 1) $\mu_+^{(j)} = 0, \mu_-^{(j)} > 0$, 2) $\mu_-^{(j)} = 0, \mu_+^{(j)} > 0$ and 3) $\mu_+^{(j)} = \mu_-^{(j)} = 0$ which correspond to the three possibilities for $\text{sign}(\lambda_j)$. For each of the d dimensions we can try all three choices. For a fixed guess of the signs let $S_+ = \{j: \lambda_j > 0\}$, $S_0 = \{j: \lambda_j = 0\}$ and $S_- = \{j: \lambda_j < 0\}$. Assuming a correct guess of $\text{sign}(\lambda_j)$ for each of the dimensions the optimization problem above reduces to the following:

PROPOSITION 2.1. *For the correct guess of $\text{sign}(\lambda_j)$ for all $j \in [d]$ it suffices to find the optimum of the above optimization problem without the constraints for $j \in S_0$. This optimum is unique.*

The proof is given in the full version. Using Proposition 2.1 and trying all guesses for $\text{sign}(\lambda_j)$ we can reduce the projection step to 3^d instances of the following optimization problem:

Minimize: $f(\mathbf{x}) = \|\mathbf{y} - \mathbf{x}\|_2^2$

Subject to: $g_i = x_i^2 - 1 \leq 0 \quad \forall i \in [n]$

$\sum_{i=1}^n w_i^{(j)} x_i = \varepsilon, \quad \forall j \in S_+; \sum_{i=1}^n w_i^{(j)} x_i = -\varepsilon, \quad \forall j \in S_-$

which can be done by finding numbers $\lambda_j > 0$ for $j \in S_+$ and $\lambda_j < 0$ for $j \in S_-$ and setting $x_i = [y_i - \sum_{j \in S_+ \cup S_-} \lambda_j w_{ij}]$. The choice of λ_j 's has to satisfy the constraints $\sum_{i=1}^n w_{ij} x_i = \varepsilon$ for all $j \in S_+$ and $\sum_{i=1}^n w_{ij} x_i = -\varepsilon$




Figure 2: One-dimensional projection. First, the initial point \mathbf{y} is moved by vector $-\lambda^* w$, which is an orthogonal vector to planes, corresponding to balance constraints. Then the resulting point is projected on the cube.

for all $j \in S_-$. In the analysis below we assume that $d = |S_+ \cup S_-|$ corresponds to the “effective dimension” of the problem.

2.3 Exact Projection Algorithms

Projection for $d = 1$. As a warm up, we first show how to perform exact projection for $d = 1$ in $\mathcal{O}(n \log n)$ time, proving Theorem 1.1 for $d = 1$. This can be further improved to $\mathcal{O}(n)$ using a more careful approach [31]. However, to the best of our knowledge, no fast algorithm is known for $d > 1$ which is the main focus of our work. Dropping the second index to simplify presentation (that is, $w_i = w_i^{(1)}$) and using the fact that $x_i = [y_i - \lambda w_i]$ we have:

$$\begin{aligned} \sum_i w_i x_i &= \sum_{i: y_i \geq 1 + \lambda w_i} w_i + \sum_{i: y_i \leq -1 + \lambda w_i} -w_i + \\ &+ \sum_{i: y_i \in (-1 + \lambda w_i, 1 + \lambda w_i)} w_i (y_i - \lambda w_i). \end{aligned}$$

We introduce notation $h_i(\lambda)$ where each h_i is the following piecewise linear function:

$$h_i(\lambda) = \begin{cases} w_i & \text{if } \lambda < (y_i - 1)/w_i \\ w_i(y_i - \lambda w_i) & \text{if } \lambda \in [(y_i - 1)/w_i, (y_i + 1)/w_i] \\ -w_i & \text{if } \lambda > (y_i + 1)/w_i \end{cases}$$

Thus $\sum_{i=1}^n w_i x_i = \sum_{i=1}^n h_i(\lambda)$ and the problem reduces to finding λ^* such that $\sum_i h_i(\lambda^*) = \pm \varepsilon$ where the sign depends on whether our dimension is in S_+ or S_- . Since $w_i \geq 0$ for all i each h_i is monotone in λ and so the function $\sum_i h_i$ is a monotone piecewise linear function. The value of λ^* can be found in $\mathcal{O}(\log n)$ iterations of binary search where each iteration requires $\mathcal{O}(n)$ time to evaluate the sum. This gives the overall running time of $\mathcal{O}(n \log n)$. See Figure 2 for an illustration.

Projection for $d = 2$. For $d = 2$ we need to find (λ_1, λ_2) such that $\sum_{i=1}^n h_i^{(j)}(\lambda_1, \lambda_2) = \pm \varepsilon$ for $j = 1, 2$, where $h_i^{(j)}(\lambda_1, \lambda_2)$




Figure 3: Two-dimensional projection. Initial point y is moved by vector $\lambda_1 w^{(1)} + \lambda_2 w^{(2)}$ and then projected on the cube.

Table 1: Theoretical properties of projection methods.

	d	Output	Time required
Alternating	any	$\mathbf{x} \in K$	Until convergence
Dijkstra's	any	projection	Until convergence
Exact (ours)	$d \leq 2$	projection	$\mathcal{O}(n \log^{d-1} n)$

is defined below.

$$h_i^{(j)}(\lambda_1, \lambda_2) = \begin{cases} w_i^{(j)} & \text{if } \sigma_i < y_i - 1 \\ w_i^{(j)}(y_i - \sigma_i) & \text{if } \sigma_i \in [y_i - 1, y_i + 1] \\ -w_i^{(j)} & \text{if } \sigma_i > y_i + 1 \end{cases}$$

where $\sigma_i = \lambda_1 w_i^{(1)} + \lambda_2 w_i^{(2)}$. The projection process is shown in Figure 3. In Appendix A, we prove Theorem 1.1 for $d = 2$ showing that nested binary search can be used to solve this problem in $\mathcal{O}(n \log n)$ time.

3. IMPLEMENTATION

3.1 Projection algorithms

We considered the following three methods for the projection step (Algorithm 1, Line 6). Their theoretical properties are summarized in Table 1.

- *Alternating projections:* A standard approach for projection on the intersection of convex sets is the alternating projections method (see [10]). It is easy to implement projections on \mathcal{B}_∞ and $\cap_{j=1}^d S_\varepsilon^j$ separately. Since both are convex bodies by alternating projections on each of them one can guarantee convergence to a point in the intersection, but there is no guarantee that this point will be the actual projection. In practice, we are able to achieve slightly better balance by modifying this approach slightly and projecting on S_0^j instead of S_ε^j . This still ensures that we get a point in the intersection in the end.
- *Dijkstra's projection:* We also consider Dijkstra's projection algorithm [1, 15]. This is a modification of the alternating projections method which is guaranteed to converge to the projection.
- *Exact projection for $d \leq 2$:* This is the algorithm presented in Section 2.2. In our experiments Dijkstra's algorithm and exact projection give similar results, since they find approximately the same projection point.

In Section 4.3 we study how quality of results produced by GD depends on choice of one of the projection methods above. Since the exact projection algorithm is computationally the

most expensive, in our experiments we mostly use the alternating projections method. Moreover, since in practice each iteration of alternating projection is computationally expensive, in the intermediate iterations we project on each plane and the cube only once, while in the last iterations we run the alternating projections method until convergence. We refer to this choice as "one-shot" alternating projection below.

3.2 Adaptive Step Size

Recall that Algorithm 1 has the following parameters: Gaussian noise variances for each step $\{\eta_t\}$ and step size parameters $\{\gamma_t\}$. Due to the spectral properties of the adjacency matrix in our experiments, the algorithm does not encounter any saddle points other than the initial point $\mathbf{x} = 0$. Therefore it suffices to add Gaussian noise only at the first iteration (that is, $\eta_t = 0$ for $t \neq 0$).

The simplest choice of the step size parameters $\{\gamma_t\}$ is constant, but it gives suboptimal results in our experiments. Carefully chosen step size parameters for different iterations not only gives better performance but can also be used to ensure that convergence can be reached in a fixed number of steps. In section 4.3 we discuss how to choose the step size to achieve good performance for a wide range of graphs.

The choice of step size parameters is complicated by the projection step. The change in the objective function and the progress towards an integral solution can both be related to the progress in Euclidean distance $\|\mathbf{x}_t - \mathbf{x}_{t+1}\|$ between the iterations, which can be greatly reduced by the projection.

Another important implementation detail is our handling of vertices which are close to integral. When the number of such vertices becomes large, the progress of the algorithm can slow down. This is due to the fact that while the gradient vector is still large, all of its large components correspond to already integral vertices and point to the outside of the feasible region. These components dominate in the computation of the projection step, which leads to slow convergence. In order to avoid this issue, we "fix" such vertices so that they become integral and no longer participate in the gradient update and the projection step. As we show in Section 4.3, this yields noticeable improvements in the quality of the results.

4. EXPERIMENTS

We design our experiments to understand how well the new partitioning algorithm behaves on real-world datasets and how it affects the performance of distributed graph processing. We are primarily interested in two quality measures of partitions: cut size (number of non-local edges) and maximum imbalance over all dimensions (maximum relative deviation of the weight in a part of the partition from the average). As pointed out in Section 1, we are not aware of an alternative scalable approach for solving the *multi-dimensional* balanced partitioning. However, some of the existing techniques for one-dimensional partitioning can be adapted for the multi-dimensional case. Next we discuss several such techniques that we evaluated.

Hash is the simplest partitioning strategy that assigns vertices to worker machines by hashing the vertex identifiers. Hashing is stateless, extremely fast in practice, and requires no preprocessing of the graph, which made it the default strategy in Giraph. The main disadvantage is that

the majority of sent messages are non-local and may result in significant communication.

Spinner is a graph partitioning algorithm that can be applied to process large-scale graphs in a distributed environment [33]. The algorithm is based on the label propagation technique in which vertices exchange their labels trying to pick the most frequent label among its neighbors. This process guarantees a high number of adjacent vertices having the same label, which are then assigned to the same worker. **Spinner** does not enforce a strict balance across partitions but integrates score functions that penalize imbalanced solutions.

BLP is another approach based on the balanced label propagation and combines the ideas of Ugander and Backstrom [42] and Meyerhenke et al. [34]. On the first step, the method creates a size-constrained clustering of the input graph using significantly more clusters than the number of available machines, k . In our implementation, we construct $c \times k$ clusters for $c = 1024$ and forbid a cluster to contain more than $\frac{|V|}{c \times k}$ vertices and $\frac{|E|}{c \times k}$ edges. On the second step, we randomly merge the clusters into k parts, which results in the multi-dimensional balance even if the original clusters have different sizes.

SHP is a distributed graph partitioner [22, 38] that is based on a classical local search heuristic [25]. Although **SHP** does not provide balancing on multiple dimensions, it supports a mode with several dimensions whose final balance is not guaranteed. The algorithm works by balancing on a new dimension, which is a combination of the specified dimensions. We configure **SHP** to find solutions having the same number of edges (with a higher coefficient in the combination) and the same number of vertices (with a lower coefficient) in every part.

We implemented the algorithms and extensively experimented with the Giraph framework, which is used as the primary tool for large-scale graph analytics at Facebook [2, 6]. Although the evaluation is performed with the single distributed graph processing system, we believe that our main conclusions are valid for other frameworks relying on the vertex-centric programming model. For our experiments, we use four large social networks that are publicly available [28]. **LiveJournal**, **Orkut**, **Twitter**, and **Friendster** are undirected graphs containing 4.8, 3.1, 41, and 65 million of vertices and 0.04, 0.12, 1.2, and 1.8 billion of edges, respectively. In addition, we experiment with several large subgraphs of the Facebook friendship graph that serve to demonstrate scalability of our approach and its performance on real-world data. We denote the graphs by **FB-X**, where **X** indicates the (approximate) number of billions of edges; this data is anonymized before processing.

Next we analyze the quality of the solutions produced by the algorithms on our dataset (Section 4.1) and evaluate various graph partitioning strategies for speeding up distributed graph processing for real-world workloads (Section 4.2). Section 4.3 investigates various parameters of **GD**.

4.1 Two-Dimensional Partitioning

Our initial experiments (see Figure 1) and earlier works [18, 29, 33] indicate that two important dimensions for the performance of Giraph jobs are the number of vertices and the number of edges. For this reason, we specify two weights for the vertices, $w_v^{(1)} = 1$ and $w_v^{(2)} = \deg(v)$ for all $v \in V$. Recall that our primary goal is to guarantee almost perfect balance for the two dimensions, as even a single overloaded

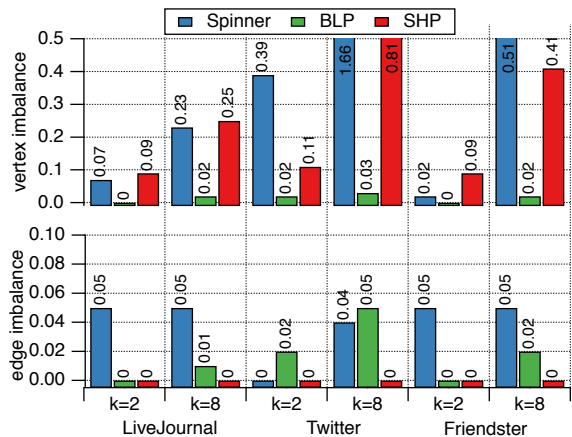


Figure 4: Vertex and edge imbalance ($\frac{\max_i w(V_i)}{\text{avg}_i w(V_i)} - 1$) of the solutions created by different algorithms on the three public networks with $k \in \{2, 8\}$. Lower values correspond to more balanced partitions. **Hash** and **GD** yield near-balanced solutions for the instances.

component affects the job performance. Figure 4 illustrates the resulting vertex and edge imbalance of the solutions on the public networks for three algorithms, **Spinner**, **BLP**, and **SHP**, using $k = 2$ and $k = 8$. The imbalance is defined as $\left(\frac{\max_i w(V_i)}{\text{avg}_i w(V_i)} - 1\right)$, where the maximum and the average are taken over the total weight of all k constructed parts. Note that this definition differs from the usual notion of imbalance in the sense that it considers only positive imbalance, since we are only interested in the most overloaded machines. We do not include the results for **Hash** and **GD**, as the corresponding values are below 0.01 for the instances.

We observe that two algorithms, **Spinner** and **SHP**, are not suitable for the multi-dimensional variant of the problem. For dense graphs with a highly skewed degree distribution (as in **Twitter**), the algorithms cannot simultaneously provide balance on the two dimensions. With the default setting, these two algorithms generate solutions in which some of the parts contain $1.5 - 2x$ more vertices than the average one. We tried to modify the techniques by adjusting relative weights of their penalty functions for vertex and degree counts. However, we were not able to design universal penalty weights that work for all instances. A similar behavior regarding the resulting balance is observed for larger graphs, **FB-3B**, **FB-80B**, and **FB-400B**. In contrast, **Hash**, **GD**, and **BLP** produced nearly-balanced (that is, having $\varepsilon \leq 0.05$ both for vertex and edge counts) solutions for all the instances. With this in mind, we exclude **Spinner** and **SHP** from further experiments.

Next we compare the quality of our algorithm as measured by the resulting edge locality, that is, the percentage of uncut edges with both endpoints in the same parts. The metric represents the fraction of local messages in Giraph jobs and corresponds to a possible reduction in communication between the worker machines. Figure 5 reports the results of **Hash**, **GD**, and **BLP** on the public dataset. Unsurprisingly, **GD** and **BLP** outperform the **Hash** algorithm in the experiment, as the latter keeps only $\frac{1}{k}$ of all the edges in the same component. The resulting edge locality of **GD** and **BLP** are close for the three graphs, though **GD** typically achieves a higher locality by 2% - 5%.




Figure 5: The percentage of local (uncut) edges produced by the three algorithms for the public graphs with $k \in \{2, 8\}$. Higher values indicate better solutions. GD achieves higher locality in all cases.




Figure 6: The percentage of local (uncut) edges produced by the three algorithms for various subgraphs of the Facebook friendship graph with $k \in \{16, 128\}$. Higher values indicate better solutions. GD achieves higher locality in all cases.

Figure 6 shows the experiments on the Facebook friendship graphs. Here we use a larger number of parts, k , which more accurately represent the real-world Giraph use case. Again, Hash produces solutions having the lowest edge localities. In fact, over 99% of the edges are cut using the partitioning strategy for an instance with a hundred partitions. This is in agreement with our measurements of the typical percentage of cross-worker Giraph messages in the production environment. On the other hand, we observe a bigger advantage of GD over BLP; the locality difference is around 10% – 20% for $k = 16$ and 5% – 10% for $k = 128$. The balanced label propagation algorithm, BLP, could be configured to produce better results by decreasing its cluster size threshold, c . However, this results in an imbalanced solution with $\epsilon > 0.05$ for the largest instance with $k = 128$. Hence, we keep the value of $c = 1024$ for all the experiments. The main difference between FB graphs and publicly available graphs is the number of edges. Therefore, we conjecture that the main reason why on FB graphs GD performs better compared to other algorithms is poor performance of existing local-search based methods on large graphs in the multi-dimensional case. This is most obvious in Figure 6 for $k = 128$ as one can see that GD is gaining a larger advantage over BLP as the size of the graph grows (3B \rightarrow 80B \rightarrow 400B). However, more experiments might be required to verify the hypothesis.

Overall we observe that GD generates solutions of higher quality than BLP and Hash on the examined instances. Therefore, we utilize the algorithm to experiment with distributed graph processing in the next section.




Figure 7: Speedup of Giraph jobs using various partitioning strategies relative to Hash measured for Page Rank (PR), Connected Components (CC), Hypergraph Clustering (HC), and Mutual Friends (MF), which are applied on FB-80B (*small*) and FB-400B (*large*) graphs. Positive values indicate improvements, negative ones indicate regressions. Vertex+edge partitioning always results in performance improvement.

4.2 Distributed Graph Processing

In this section we conduct an experimental evaluation of various graph partitioning strategies for speeding up distributed graph processing. Here we argue and experimentally demonstrate that multi-dimensional balancing is a suitable objective for the application. We experiment with four graph algorithms implemented in Giraph. **Page Rank** and **Connected Components**, are popular benchmarks for verifying the performance of distributed systems. **Page Rank** iteratively propagates vertex ranks through adjacent edges; our implementation performs 30 iterations for the algorithm. For the **Connected Components** algorithm, we use a simple label propagation technique in which vertices iteratively update their labels based on the minimum label of their neighbors; for our graphs, the process converges after at most 50 rounds. The other two algorithms, **Hypergraph Clustering** and **Mutual Friends**, are applications for large-scale graph analytics. The former is used to find a certain clustering of the input graph by converting it to a hypergraph. The latter builds a set of features for friend recommendation. Both applications extensively exchange messages between adjacent vertices, which adds a significant communication overhead.

Figure 7 depicts the results of our experiment. Since we are interested in the impact of various partitioning policies on the performance of Giraph, we report the relative differences to the baseline policy, Hash. Here we measure the total runtime of an application using GD as the partitioning strategy in three modes, *vertex partitioning* (one-dimensional balance on vertex count), *edge partitioning* (balance on edge count), and *vertex-edge partitioning* (two-dimensional balance both on vertex and edge counts). Every algorithm is applied in two configurations, *small* and *large*. The first one uses the FB-80B graph and a cluster with 16 worker machines, while the second one process FB-400B using 128 workers.

The key finding is that one-dimensional partitioning cannot provide consistent benefits across all the Giraph applications. In fact, we observe performance regression for some instances, in particular, when the number of utilized worker machines is large, that is, $k = 128$. In this scenario, we notice a few workers whose running time is significantly larger than the average; see Figure 1. Since in Giraph (and other vertex-

Table 2: Impact of partitioning policy on the running time and the amount of sent messages across 128 Giraph workers for the **Page Rank** application applied on the FB-400B graph. The numbers are average values over 30 iterations.

Partitioning	Runtime, sec			Communication, GB		
	mean	max	stdev	mean	max	stdev
Hash	95	102	27	69.5	69.6	2.4
vertex	93	143	25	18.6	47.6	6.8
edge	82	120	22	25.7	38.2	5.9
vertex-edge	84	88	21	29.1	30.6	2.8

centric systems) the computation is split into a number of supersteps that end with a global synchronization barrier, the performance is determined by the slowest worker. Notice that a similar phenomena regarding the *vertex partitioning* has been observed in earlier works [3, 18, 20, 40]. In contrast, the two-dimensional partitioning always results in a speedup over the default **Hash** strategy. The improvement is in the order of 10% – 30% for the examined applications.

To get a deeper understanding of the source of performance differences, we analyze the detailed logs for the **Page Rank** application using a cluster with 128 worker machines. Table 2 shows the measurements of the mean, maximum, and standard deviation of the time to compute a superstep by all the workers. The results indicate that the with the hash-based partitioning, the workers are idling on average for 7 seconds per superstep waiting for the slowest one to complete the work. With one-dimensional partitioning the idling time is much longer, 50 seconds for vertex-based partitioning and 38 seconds for edge-based one, which is the primary reason for the performance regression. The two-dimensional partitioning results in a more even load across the workers delivering a 13.2% speedup. Table 2 also indicates a significant communication reduction over the baseline partitioning, as measured by the total size of messages sent between the workers via network. For the **Page Rank** application, the average reduction is correlated with the edge locality of the corresponding partitioning. However, an unbalanced partitioning causes some workers to use more memory resources and become a bottleneck for graph processing.

Finally, we emphasize that the timings analyzed in the section exclude the running times of the partitioner itself. This is realistic for our use case in which the same friendship graph is expected to be processed multiple times for various tasks. Thus, the extra overhead incurred by a partitioning strategy is amortized among several runs.

4.3 Parameters of GD

In this section we perform an experimental comparison of various choices of the projection step algorithm in GD and study its convergence properties. Unless specified otherwise, we use two-dimensional GD in the following setting: 1) balance is required with respect to the number of vertices and their degrees, 2) in the projection step we use “one-shot” alternating projection (see Section 3.1), 3) we use adaptive step size and vertex fixing as described in Section 3.2.

Since behavior of gradient descent algorithms can depend on selection of the step size parameters, we used experiments to establish convergence of GD with different choices of these parameters. In particular, our implementation aims to ensure that the step length $\|\mathbf{x}^{(t)} - \mathbf{x}^{(t+1)}\|_2$ remains close to constant




Figure 8: Comparison of step choices for GD with fixed step length, that is, $\|\mathbf{x}^{(t)} - \mathbf{x}^{(t+1)}\|_2 = const$, for 100 iterations and $\xi = \sqrt{n}/100$. Step length $2 \cdot \xi$ results in good performance.

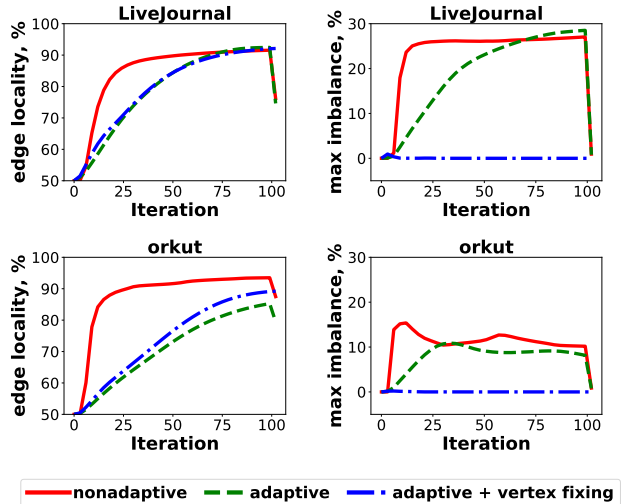


Figure 9: Quality and imbalance comparison of GD 1) without adaptive step size, 2) with adaptive step size and 3) with adaptive step size with vertex fixing. The left side shows edge locality and the right side – maximum imbalance over all dimensions. For nonadaptive and adaptive strategies the changes in the number of cut edges and imbalance in the last iteration are due to fixing in the end of the algorithm the accumulated imbalance resulting from “one-shot” alternating projection. Using GD with adaptive step size and vertex fixing results in better locality and preserves almost perfect balance during algorithm execution.

between iterations. A natural scaling parameter for the step length is \sqrt{n} as it corresponds to the distance between the initial solution $\mathbf{x}_0 = 0$ and an integral solution of the form $\{-1, 1\}^n$. As we show in Figure 8 for various graphs, a good choice of step size is $2\frac{\sqrt{n}}{100}$, where 100 is the limit we set on the number of iterations due to the constraints on the execution running time.

In Figure 9 we show how adaptive step size and vertex fixing affect the performance of the algorithm. Note that compared with other methods vertex fixing not only improves quality but also preserves almost perfect balance even when simple “one-shot” alternating projection is used. Finally, in Figure 10 we show analysis of performance of the algorithm under different choices of the projection method. The results show that the exact projection algorithm with sufficiently large allowed imbalance leads to the best perfor-




Figure 10: Quality comparison of GD with various projection methods. We compare exact projection with various allowed imbalance parameters and “one-shot” alternating projection. Allowing more imbalance typically results in higher quality. “One-shot” alternating projection, which we choose as our default implementation option due to its efficiency, produces results comparable with the exact projection. Dykstra’s projection produces the same results as the exact projection, and therefore is not shown.

mance. Larger imbalance permits more partitions, possibly including ones with better locality, allowing the overall algorithm to find results with higher locality. However, the alternating projections algorithm can often be used to achieve similar performance. This is likely due to the fact that the alternating projections algorithm despite not computing the projection outputs a point close enough to it.

4.4 Multi-Dimensional Experiments

We perform experiments for $d = 3$ and $d = 4$ to illustrate the performance of our algorithms in the multi-dimensional case. For these multidimensional experiments in addition to balancing on the number of vertices and edges we also balance based on the following additional vertex weights:

- *Pagerank*. We use Pagerank to model activity level of a node. High Pagerank likely means that the vertex is accessed often; hence, balancing on the value can be beneficial for load balancing purposes.
- *Sum of neighbor degrees*. We also use the sum of degrees over neighbors of a vertex as a weight function. We choose the sum of neighbor degrees as a proxy for the size of the 2-hop neighborhood of a vertex, which is computationally expensive to compute for large graphs.

We run experiments on `LiveJournal` and `Orkut` graphs, as well as on `sx-stackoverflow` graph (2.6M vertices, 28M edges) – the largest SNAP graph which is not a social network. We compare our algorithm with `METIS` with allowed imbalance of 0.5%. The results are presented in Table 3. They indicate that `METIS` achieves poor balance for multiple constraints and that GD outperforms `METIS` by almost all parameters in most cases.

4.5 Performance Analysis

Finally, we analyze scalability of our algorithm. Our results are obtained on a Hadoop cluster of 128 workers; each of the machines is a dual-node 2.4 GHz Intel Xeon E5-2680 with 256GB RAM. Figure 11 reports the running time of GD in machine-hours on `FB-X` graphs of various size with balance on two dimensions. We observe a near-linear growth of the running time with the size of the input graph. In comparison,

the running time of the SHP algorithm exceed the values by a factor of 1.5 – 2 on the same cluster configuration.

5. CONCLUSION

We introduced a new algorithm for the MULTI-DIMENSIONAL BALANCED GRAPH PARTITIONING problem, which produces balanced partitions according to multiple user-specified weight functions while maintaining high edge locality. Our computational experiments indicate that this algorithm is scalable for large real-world graphs, and outperforms existing techniques for the problem. Resulting partitions allow one to achieve moderate speedups in computational time for several real-world Giraph workloads. This is in contrast with balancing on just one dimension, that can result in worsen performance.

One of the interesting directions for future work is incorporating a wider range of balancing requirements, for example, those that can depend on the resulting partitioning itself, such as the number of local edges and the maximum number of edges going between any pair of parts in the resulting partition. The latter quantity can substantially affect performance of the distributed computation in Giraph-like systems, since communication between different machines depends on the number of edges going between the corresponding parts. The proposed algorithm cannot directly handle such solution-dependent weight functions as they are not expressible through an a priori fixed collection of weight functions.

An interesting theoretical question is finding a fast algorithm for exact projection for $d > 2$. As we show in the full version, it is possible to use nested binary search to find $\{\lambda_j\}$ (and therefore the projection) with arbitrary precision. Unfortunately, the running time of the suggested algorithm is unknown, because it is unclear how to estimate left and right bounds for the binary search. Determining these bounds gives an algorithm with running time $\mathcal{O}(n \cdot \prod_{i=1}^d \log \frac{r_j - l_j}{\delta})$, where l_j and r_j are bounds for λ_j and δ is the required precision.

Another interesting theoretical question is understanding the convergence properties of our algorithm (or a similar gradient descent based method) under some assumption about the spectral properties of the graph. We see this as a challenging open problem. While noisy gradient descent is known to have fast convergence for non-convex optimization subject to equality constraints, if inequality constraints are allowed, convergence analysis is open [16].

APPENDIX

A. PROJECTION FOR $D = 2$

In this section we introduce a randomized $O(n \log n)$ -time algorithm for finding projection for $d = 2$. Recall from Section 2.2 that for $\mathbf{y} \in \mathbb{R}^n$ we need to find $\boldsymbol{\lambda}^* = (\lambda_1^*, \lambda_2^*)$ such that $h^{(1)}(\boldsymbol{\lambda}^*) = c_1$ and $h^{(2)}(\boldsymbol{\lambda}^*) = c_2$. For $\boldsymbol{\lambda} = (\lambda_1, \lambda_2)$ we define $h^{(j)}(\boldsymbol{\lambda}) = \sum_{i=1}^n h_i^{(j)}(\boldsymbol{\lambda})$ for $j \in \{1, 2\}$, where

$$h_i^{(j)}(\boldsymbol{\lambda}) = \begin{cases} w_i^{(j)} & \text{if } \sum_k \lambda_k w_i^{(k)} < y_i - 1 \\ -w_i^{(j)} & \text{if } \sum_k \lambda_k w_i^{(k)} > y_i + 1 \\ w_i^{(j)}(y_i - \sum_k \lambda_k w_i^{(k)}) & \text{otherwise.} \end{cases}$$

Once we find $(\lambda_1^*, \lambda_2^*)$ we can compute the coordinates of \mathbf{x} as $x_i = [y_i - w_i^{(1)} \lambda_1^* - w_i^{(2)} \lambda_2^*]$. We introduce an auxiliary

Table 3: Comparison of GD with METIS. The results show that for high-dimensional balanced partitioning, METIS cannot guarantee balance across all dimensions. Better results shown in bold.

		LiveJournal		Orkut		sx-stackoverflow	
		GD	METIS	GD	METIS	GD	METIS
$d = 2$: balance on vertices and degrees	locality, %	91.71	93.74	88.36	86.52	75.82	80.41
	imbalance	0.04	0.5	0.02	0.7	0.04	0.6
	memory, MB	2635	4085	4673	10259	1587	4113
	time, sec	117	44	203	92	68	55
$d = 3$: balance on vertices, degrees and sum of neighbor degrees	locality, %	88.74	73.36	89.55	62.1	76.8	60.09
	imbalance	0.05	30	0.02	1.6	0.1	6.5
	memory, MB	2711	4802	4697	12271	1627	4985
	time, sec	140	66	196	303	76	131
$d = 4$: balance on vertices, degrees, sum of neighbor degrees and pagerank	locality, %	87.93	74.36	75.58	65.08	77.04	78.54
	imbalance	0.5	38	2.7	20	0.4	3.8
	memory, MB	2939	4839	4896	12294	1754	5013
	time, sec	227	66	240	297	88	142




Figure 11: Scalability of the distributed implementation of GD on FB-X graphs of various size. The results indicate a near-linear dependence of the running time on the number of edges in the graph.

function Δ which we use to solve the above problem using binary search:¹

DEFINITION A.1. *Suppose that λ_1 is such that there exists λ_2 for which the constraint $h^{(2)}(\lambda_1, \lambda_2) = c_2$ is satisfied. Then we define $\Delta(\lambda_1) \triangleq h^{(1)}(\lambda_1, \lambda_2)$.*

We now describe an $\mathcal{O}(n \log n)$ -time algorithm for finding $(\lambda_1^*, \lambda_2^*)$. The algorithm is shown as Algorithm 2. It takes as a parameter a Boolean variable Δ^+ indicating whether Δ is an increasing or decreasing function. We run the algorithm under both assumptions and select a solution satisfying the constraints.

We outline the main ideas behind Algorithm 2 below. Consider the (λ_1, λ_2) plane partitioned by the following lines (which we call *boundary lines*):

$$\begin{aligned} y_i - \lambda_1 w_i^{(1)} - \lambda_2 w_i^{(2)} &= 1 \\ y_i - \lambda_1 w_i^{(1)} - \lambda_2 w_i^{(2)} &= -1, \end{aligned}$$

for all i . Let L be the set of boundary lines (line 2). We refer to the subsets of the plane resulting from its partition by the boundary lines as *regions* (see Figure 12 where the regions are referred to as $\{T_i\}$). Boundary lines separate the plane into half-planes corresponding to the different cases in the definitions of the corresponding $h_i^{(j)}$. Therefore, inside each region all $h_i^{(j)}$ are linear and hence $h^{(j)}$ are also linear.

¹In the full version of the paper we will show that the Δ is well-defined and monotone.

The intuition behind the algorithm is then as follows (in order to achieve the best performance the exact details differ slightly from this simplified presentation). Suppose we could find a region that contains some solution λ^* . Then since constraint functions are linear inside the region, in order to find λ^* we could solve a system of linear equations over λ_1 and λ_2 . We identify such region, with binary search over λ_1 by using monotonicity of Δ . We consider only a finite set of values: λ_1 -coordinates of intersections of boundary lines. Since there are $\mathcal{O}(n)$ boundaries, there are $\mathcal{O}(n^2)$ intersections (e.g., in Figure 12 we consider only points a, b, c and d). Hence $\mathcal{O}(\log n)$ iterations of binary search suffice. The only difference between Algorithm 2 and the above approach is that after the binary search on λ_1 we still have to try $\mathcal{O}(n)$ regions to identify the exact region which contains λ^* (see Algorithm 2 for the details).

Now consider one iteration of the binary search. Let λ_1^l and λ_1^r be its current boundaries. Let Λ be a set of all intersection points (λ_1, λ_2) such that $\lambda_1 \in (\lambda_1^l, \lambda_1^r)$. Since Δ is monotone, for any λ_1^l we can use binary search by checking whether λ^* is greater or less than λ_1^l through a comparison of $\Delta(\lambda_1^l)$ and c_1 (lines 9-12). Computing $\Delta(\lambda_1^l)$ requires solving the one-dimensional problem over λ_2 discussed in Section 2.3 and thus can be done in $\mathcal{O}(n)$ time.

In order to have binary search run in $\mathcal{O}(\log n)$ iterations it suffices to find a value $\lambda_1^l \in (\lambda_1^l, \lambda_1^r)$ which with constant probability splits Λ into two subsets of points, those with $\lambda_1 > \lambda_1^l$ and with $\lambda_1 < \lambda_1^l$ respectively, of size at most $\frac{2}{3}n$ each. In particular, it suffices to sample a uniformly random point $(\lambda_1^l, \lambda_2^l)$ from Λ . The following lemma bounds the overall running time of these sampling steps.

LEMMA A.1. *The overall time required for sampling random points from Λ in line 8 of Algorithm 2 is $\mathcal{O}(n \log n)$.*

PROOF. Consider three cases:

Case 1. $|\Lambda| > n \log n$. We sample $\mathcal{O}(n)$ uniformly random pairs of lines from L and find an intersection of each pair (assume no parallel lines which can be handled separately). Since the number of lines is $\mathcal{O}(n)$ w.h.p. we sample at least one intersection which lies in Λ . The last condition can be checked in $\mathcal{O}(n)$ time and if it doesn't hold then we conclude that w.h.p. $|\Lambda| \leq n \log n$. We then compute S , the set of all points in Λ in $\mathcal{O}(n \log n)$ time and proceed to the second case. We show how to compute Λ in the full version.

Algorithm 2: Function returning λ_1^*, λ_2^* for given 2-dimensional problem.

```

input :  $\{w_i^{(1)}\}, \{w_i^{(2)} > 0\}, \{y_i\}, c_1, c_2$ 
parameter:  $\Delta^+ \in \{true, false\}$  indicating whether  $\Delta$ 
increases
output :  $(\lambda_1^*, \lambda_2^*)$ 
1 Function Project-2D
   /*  $L =$  set of lines parameterized by
    $(y_i, w_i^{(1)}, w_i^{(2)}, \pm 1)$  corresponding to lines of
   the form  $y_i - \lambda_1 w_i^{(1)} - \lambda_2 w_i^{(2)} = \pm 1$  */
2  $L := \{(y_i, w_i^{(1)}, w_i^{(2)}, \pm 1) \mid i \in [n]\}$ 
3  $\lambda_1^l := -\infty, \lambda_1^r := +\infty$ 

   /* Run binary search */
4 while true do
5    $\Lambda :=$  set of intersection points  $(\lambda_1, \lambda_r)$  of
   lines in  $L$  such that  $\lambda_1 \in (\lambda_1^l, \lambda_1^r)$ 
6   if  $\Lambda = \emptyset$  then
7     break
8   Sample a uniformly random intersection
   point  $(\lambda_1', \lambda_2')$  from  $\Lambda$ 
9   if  $\Delta(\lambda_1') > c_1$  then
10    | If  $\Delta^+$  set  $\lambda_1^r := \lambda_1'$ , otherwise set  $\lambda_1^l := \lambda_1'$ 
11   else
12    | If  $\Delta^+$  set  $\lambda_1^l := \lambda_1'$ , otherwise set  $\lambda_1^r := \lambda_1'$ 

13 Let  $\{R_t\}_{t=1}^T$  be a partition of  $(\lambda_1^l, \lambda_1^r) \times \mathbb{R}$  by
   boundary lines, sorted from bottom to top
14 Compute coefficients for the system of
   linear equations for  $R_1$  (as in Theorem A.3)
15 for  $t = 1 \dots T$  do
16   Let  $(\lambda_1, \lambda_2)$  be a solution to the linear
   system for  $R_t$ 
17   if  $(\lambda_1, \lambda_2) \in R_t$  then
18     return  $(\lambda_1, \lambda_2)$ 
19   Update the coefficients corresponding to
   crossing the boundary line between  $R_t$ 
   and  $R_{t+1}$  as shown in Theorem A.3

```

Case 2. $n \leq |\Lambda| \leq n \log n$. Note that in this case $\Lambda = \{(\lambda_1, \lambda_2) \in S \mid \lambda_1 \in (\lambda_1^l, \lambda_1^r)\}$, where S is as defined above. We sample $\mathcal{O}(n)$ random points from S so that w.h.p. we get at least one point from Λ . As before, if this doesn't happen, we conclude that w.h.p. $|\Lambda| < n$ and proceed to the last case.

Case 3. $|\Lambda| < n$. In this case we maintain Λ directly. When we sample a random point $(\lambda_1', \lambda_2') \in \Lambda$, we remove from Λ all points on one of the sides from λ_1' as directed by the binary search.

In each of the cases above each of $\mathcal{O}(\log n)$ iterations takes $\mathcal{O}(n)$ time and pre-/post-processing between the cases takes $\mathcal{O}(n \log n)$ time. Thus sampling takes $\mathcal{O}(n \log n)$ time. ■

Using the above algorithm we can find λ_1^l and λ_1^r such that there are no intersection points between them. Since there are $\mathcal{O}(\log n)$ iterations and each of them requires $\mathcal{O}(n)$ time on average, the total running time is $\mathcal{O}(n \log n)$. This completes a proof of the following theorem (corresponding to lines 2-12 of the algorithm).




Figure 12: Example of regions for $n = 2, \mathbf{y} = (0, 0)$. Boundary lines are $\lambda_1 + \lambda_2 = \pm 1$ and $\lambda_2 = \pm 1$. These lines partition (λ_1, λ_2) -plane into nine regions T_1, T_2, \dots, T_9 . Intersection points are a, b, c and d . Current intersection points considered by the algorithm (between λ_1^l and λ_1^r) are shown in red.

THEOREM A.2. *There exists an $\mathcal{O}(n \log n)$ -time randomized algorithm returning λ_1^l and λ_1^r such that:*

1. *No intersections of boundary lines in $[\lambda_1^l, \lambda_1^r]$,*
2. *There exists a solution $(\lambda_1^\dagger, \lambda_2^\dagger)$ such that $\lambda_1^\dagger \in [\lambda_1^l, \lambda_1^r]$.*

After we find λ_1^l and λ_1^r as in Theorem A.2 we show that there are only $\mathcal{O}(n)$ regions which can contain a solution and we can check them in $\mathcal{O}(n \log n)$ time. The following theorem completes the proof of Theorem 1.1 for $d = 2$:

THEOREM A.3. *If there exists a solution λ^\dagger such that $\lambda_1^\dagger \in (\lambda_1^l, \lambda_1^r)$ and no intersection points are between $(\lambda_1^l, \lambda_1^r)$ then λ^* can be found in $\mathcal{O}(n \log n)$ time.*

PROOF. We show how to find λ^* in lines 13-19 of the algorithm. Consider set $S = (\lambda_1^l, \lambda_1^r) \times \mathbb{R}$. Let $\{R_t\}_{t=1}^T$ be the partition of S into parts lying between the boundary lines. Since S doesn't contain boundary intersections and there are $\mathcal{O}(n)$ boundaries, the size of the partition is $\mathcal{O}(n)$. For each R_t we solve the following system of equations over λ_1 and λ_2 :

$$\begin{cases} \sum_{i=1}^n h_i^{(1)}(\lambda_1, \lambda_2) = c_1, \\ \sum_{i=1}^n h_i^{(2)}(\lambda_1, \lambda_2) = c_2 \end{cases}$$

Since no boundary line crosses R_t , it is a subset of some region. Therefore, $h_i^{(1)}$ and $h_i^{(2)}$ are linear inside R_t , meaning that the above system becomes a system of linear equations. If the solution to the system belongs to R_t , then we can take it as λ^* . Thus it only remains to show how to find coefficients for the system in $\mathcal{O}(n \log n)$ total time.

Recall that in Algorithm 2 we assume that $\{R_t\}$ are sorted from bottom to top. For R_1 we find the linear system coefficients in $\mathcal{O}(n)$ time. Assume that the R_t are already computed. To find the coefficients for next set R_{t+1} , notice that R_t and R_{t+1} are separated by some boundary line. This line corresponds to some $h_i^{(j)}$ and therefore crossing it will change the coefficient of only this $h_i^{(j)}$, and the coefficients can be recomputed in $\mathcal{O}(1)$ time. Since there are $\mathcal{O}(n)$ boundary lines, the overall time for recomputation is also $\mathcal{O}(n)$. Taking sorting of $\{R_t\}$ into account, the total running time is $\mathcal{O}(n \log n)$. ■

B. REFERENCES

- [1] Dykstra's projection algorithm. https://en.wikipedia.org/wiki/Dykstra%27s_projection_algorithm. Accessed: 2019-02-13.
- [2] Apache Giraph. <http://giraph.apache.org/>.
- [3] Z. Abbas, V. Kalavri, P. Carbone, and V. Vlassov. Streaming graph partitioning: An experimental study. *PVLDB*, 11(11):1590–1603, 2018.
- [4] A. Amir, J. Ficler, R. Krauthgamer, L. Roditty, and O. S. Shalom. Multiply balanced k -partitioning. In *LATIN 2014: Theoretical Informatics - 11th Latin American Symposium, Montevideo, Uruguay, March 31 - April 4, 2014. Proceedings*, pages 586–597, 2014.
- [5] A. Anandkumar and R. Ge. Efficient approaches for escaping higher order saddle points in non-convex optimization. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, pages 81–102, 2016.
- [6] C. Avery. Giraph: Large-scale graph processing infrastructure on Hadoop. *Proceedings of the Hadoop Summit. Santa Clara*, 11(3):5–9, 2011.
- [7] K. Aydin, M. Bateni, and V. S. Mirrokni. Distributed balanced partitioning via linear embedding. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, San Francisco, CA, USA, February 22-25, 2016*, pages 387–396, 2016.
- [8] D. P. Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.
- [9] C.-E. Bichot and P. Siarry. *Graph partitioning*. John Wiley & Sons, 2013.
- [10] S. Boyd and J. Dattorro. Alternating projections. 2003.
- [11] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [12] A. Buluç, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent advances in graph partitioning. In *Algorithm Engineering - Selected Results and Surveys*, pages 117–158. 2016.
- [13] D. Delling, A. V. Goldberg, I. P. Razenshteyn, and R. F. F. Werneck. Exact combinatorial branch-and-bound for graph bisection. In *Proceedings of the 14th Meeting on Algorithm Engineering & Experiments, ALENEX 2012, The Westin Miyako, Kyoto, Japan, January 16, 2012*, pages 30–44, 2012.
- [14] L. Dhulipala, I. Kabiljo, B. Karrer, G. Ottaviano, S. Pupyrev, and A. Shalita. Compressing graphs and indexes with recursive graph bisection. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, pages 1535–1544, New York, NY, USA, 2016. ACM.
- [15] R. L. Dykstra. An algorithm for restricted least squares regression. *Journal of the American Statistical Association*, 78(384):837–842, 1983.
- [16] R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points - online stochastic gradient for tensor decomposition. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 797–842, 2015.
- [17] R. Ge, J. D. Lee, and T. Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2973–2981, 2016.
- [18] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin. Powergraph: distributed graph-parallel computation on natural graphs. In *OSDI*, volume 12, page 2, 2012.
- [19] J. E. Gonzalez, R. S. Xin, A. Dave, D. Crankshaw, M. J. Franklin, and I. Stoica. GraphX: graph processing in a distributed dataflow framework. In *Proceedings of the 11th USENIX conference on Operating Systems Design and Implementation*, pages 599–613. USENIX Association, 2014.
- [20] Y. Guo, S. Hong, H. Chafi, A. Iosup, and D. Epema. Modeling, analysis, and experimental comparison of streaming graph-partitioning policies. *Journal of Parallel and Distributed Computing*, 108:106–121, 2017.
- [21] P. Jain and P. Kar. Non-convex Optimization for Machine Learning. *ArXiv e-prints*, Dec. 2017.
- [22] I. Kabiljo, B. Karrer, M. Pundir, S. Pupyrev, A. Shalita, Y. Akhremtsev, and A. Presta. Social hash partitioner: A scalable distributed hypergraph partitioner. *PVLDB*, 10(11):1418–1429, 2017.
- [23] G. Karypis and V. Kumar. Metis – unstructured graph partitioning and sparse matrix ordering system, version 2.0. Technical report, 1995.
- [24] G. Karypis and V. Kumar. Multilevel algorithms for multi-constraint graph partitioning. In *Proceedings of the 1998 ACM/IEEE Conference on Supercomputing, SC '98*, pages 1–13, Washington, DC, USA, 1998. IEEE Computer Society.
- [25] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell system technical journal*, 49(2):291–307, 1970.
- [26] R. Krauthgamer, J. Naor, and R. Schwartz. Partitioning graphs into balanced components. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 942–949, 2009.
- [27] J. D. Lee, B. Recht, R. Salakhutdinov, N. Srebro, and J. A. Tropp. Practical large-scale optimization for max-norm regularization. In *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada.*, pages 1297–1305, 2010.
- [28] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [29] L. Li, R. Geda, A. B. Hayes, Y. Chen, P. Chaudhari, E. Z. Zhang, and M. Szegedy. A simple yet effective balanced edge partition model for parallel computing. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(1):14, 2017.
- [30] Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein. Distributed graphlab: a framework for machine learning and data mining in the cloud. *PVLDB*, 5(8):716–727, 2012.
- [31] N. Maculan, C. P. Santiago, E. Macambira, and M. Jardim. An $O(n)$ algorithm for projecting a vector on the intersection of a hyperplane and a box in \mathbb{R}^n .

- Journal of optimization theory and applications*, 117(3):553–574, 2003.
- [32] K. Makarychev and Y. Makarychev. Nonuniform graph partitioning with unrelated weights. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pages 812–822, 2014.
- [33] C. Martella, D. Logothetis, A. Loukas, and G. Siganos. Spinner: Scalable graph partitioning in the cloud. In *33rd IEEE International Conference on Data Engineering, ICDE 2017, San Diego, CA, USA, April 19-22, 2017*, pages 1083–1094, 2017.
- [34] H. Meyerhenke, P. Sanders, and C. Schulz. Partitioning complex networks via size-constrained clustering. In *International Symposium on Experimental Algorithms*, pages 351–363. Springer-Verlag New York, Inc., 2014.
- [35] J. Nishimura and J. Ugander. Restreaming graph partitioning: Simple versatile algorithms for advanced balancing. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1106–1114, New York, NY, USA, 2013. ACM.
- [36] S. Ou, K. Yang, and A. Liotta. An adaptive multi-constraint partitioning algorithm for offloading in pervasive systems. In *4th IEEE International Conference on Pervasive Computing and Communications (PerCom 2006), 13-17 March 2006, Pisa, Italy*, pages 116–125, 2006.
- [37] K. Schloegel, G. Karypis, and V. Kumar. A new algorithm for multi-objective graph partitioning. In *Euro-Par '99 Parallel Processing, 5th International Euro-Par Conference, Toulouse, France, August 31 - September 3, 1999, Proceedings*, pages 322–331, 1999.
- [38] A. Shalita, B. Karrer, I. Kabiljo, A. Sharma, A. Presta, A. Adcock, H. Killapi, and M. Stumm. Social Hash: an assignment framework for optimizing distributed systems operations on social networks. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, pages 455–468. USENIX Association, 2016.
- [39] J. Sun, Q. Qu, and J. Wright. When are nonconvex problems not scary? *CoRR*, abs/1510.06096, 2015.
- [40] J. Sun, H. Vandierendonck, and D. S. Nikolopoulos. VEBO: A vertex-and edge-balanced ordering heuristic to load balance parallel graph processing. *arXiv preprint arXiv:1806.06576*, 2018.
- [41] C. E. Tsourakakis, C. Gkantsidis, B. Radunovic, and M. Vojnovic. FENNEL: streaming graph partitioning for massive scale graphs. In *Seventh ACM International Conference on Web Search and Data Mining, WSDM 2014, New York, NY, USA, February 24-28, 2014*, pages 333–342, 2014.
- [42] J. Ugander and L. Backstrom. Balanced label propagation for partitioning massive graphs. In *Sixth ACM International Conference on Web Search and Data Mining, WSDM 2013, Rome, Italy, February 4-8, 2013*, pages 507–516, 2013.
- [43] S. Verma, L. M. Leslie, Y. Shin, and I. Gupta. An experimental comparison of partitioning strategies in distributed graph processing. *PVLDB*, 10(5):493–504, 2017.
- [44] S. Wright and J. Nocedal. Numerical optimization. *Springer Science*, 35(67-68):7, 1999.